Pearson information-based lower bound on Fisher information

Dave Zachariah and Petre Stoica

arXiv:1611.07712v1 [math.ST] 23 Nov 2016

***Abstract*—The Fisher information matrix (FIM) plays an important role in the analysis of parameter inference and system design problems. In a number of cases, however, the statistical data distribution and its associated information matrix are either unknown or intractable. For this reason, it is of interest to develop useful lower bounds on the FIM. In this lecture note, we derive such a bound based on moment constraints. We call this bound the Pearson information matrix (PIM) and relate it to properties of a misspecified data distribution. Finally, we show that the inverse PIM coincides with the asymptotic covariance matrix of the optimally weighted generalized method of moments.**

The Cramér-Rao bound (CRB) is a useful tool for the analysis of parameter inference problems, benchmarking estimators, and system design [1]–[4]. Let $\mathbf{y} = [y_1 \cdots y_N]^\top$ denote the observed data from a system and let $\boldsymbol{\theta} \in \mathbb{R}^n$ be the parameters of interest. The CRB exists under certain regularity conditions and is given by the inverse of the Fisher information matrix (FIM) $\mathbf{J}(\boldsymbol{\theta})$, which is a function of the probability density $p(\mathbf{y};\boldsymbol{\theta})$. More specifically, the FIM is defined as

$$\mathbf{J}(\boldsymbol{\theta}) \triangleq \mathrm{E}\left[\frac{\partial \ln p(\mathbf{y};\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \frac{\partial \ln p(\mathbf{y};\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}^\top\right] \succeq \mathbf{0}, \tag{1}$$

where the gradient $\frac{\partial \ln p(\mathbf{y};\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}$ is known as the 'score function'. Under certain regularity conditions, the score function has zero mean

$$\mathrm{E}\left[\frac{\partial \ln p(\mathbf{y};\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right] = \mathbf{0}. \tag{2}$$

In many applications, $\mathbf{J}(\boldsymbol{\theta})$ may not be obtainable. For instance, $p(\mathbf{y};\boldsymbol{\theta})$ may be unknown in a practical problem. As an example, let the observations be modeled as $\mathbf{y} = f(\mathbf{u})$, where $f(\cdot)$ is a nonlinear function and $\mathbf{u}$ follows a probability density function $p(\mathbf{u};\boldsymbol{\theta})$. Then a closed-form expression for $p(\mathbf{y};\boldsymbol{\theta})$ is not available in general even if $p(\mathbf{u};\boldsymbol{\theta})$ is known.

It is possible, however, to derive tractable lower bounds on the FIM. If $\mathrm{E}[\mathbf{y}]$ is a function of $\boldsymbol{\theta}$ and $\mathrm{Cov}[\mathbf{y}]$ is independent of $\boldsymbol{\theta}$, then using the Gaussian distribution in lieu of $p(\mathbf{y};\boldsymbol{\theta})$ leads to the minimum FIM $\mathbf{J}_G(\boldsymbol{\theta}) \preceq \mathbf{J}(\boldsymbol{\theta})$ (where G stands for Gaussian) and therefore the 'worst-case' inference scenario [5], [6]. A generalization to $\boldsymbol{\theta}$-dependent $\mathrm{Cov}[\mathbf{y}]$ was given in [7]. The minimum FIM can be used for robust system design and estimator formulations. However, in some cases the minimum information can be overly conservative.

In this lecture note, we derive a tighter lower bound

$$\mathbf{0} \preceq \mathbf{L}(\boldsymbol{\theta}) \preceq \mathbf{J}(\boldsymbol{\theta}) \in \mathbb{R}^{n\times n}, \tag{3}$$

based on moment constraints. For this reason we call $\mathbf{L}(\boldsymbol{\theta})$ the Pearson information matrix (PIM) as an homage to the inventor of 'the method of moments' [8]. As we will see, the PIM is related to the generalized method of moments [9].

This work has been partly supported by the Swedish Research Council (VR) under contract 621-2014-5874.

I. Relevance

The PIM is a tractable tool for analyzing parameter estimation and system design problems when the statistical data distribution is unknown or intractable.

II. Prerequisities

The reader needs basic knowledge about linear algebra, elementary probability theory, and statistical signal processing.

III. Preliminaries

We begin by constructing a function $\mathbf{z}(\mathbf{y})$ that contains M statistics of $\mathbf{y}$. We assume that $\mathbf{z}(\mathbf{y})$ has computable—either analytically or numerically—mean and covariance

$$\begin{aligned}\boldsymbol{\mu}(\boldsymbol{\theta}) &\triangleq \mathrm{E}[\mathbf{z}(\mathbf{y})] \in \mathbb{R}^M,\\ \boldsymbol{\Sigma}(\boldsymbol{\theta}) &\triangleq \mathrm{E}[(\mathbf{z}(\mathbf{y}) - \boldsymbol{\mu})\,(\mathbf{z}(\mathbf{y}) - \boldsymbol{\mu})^\top] \in \mathbb{R}^{M\times M},\end{aligned} \tag{4}$$

where $M \geq n$. For instance, $\mathbf{z}$ may be constructed using powers of the data, that is, its elements are made up of empirical moments $\{y_i\}$ $\{y_i y_j\}$, $\{y_i y_j y_k\}$, etc. We assume that $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ is nonsingular. For notational simplicity, we drop the argument $\boldsymbol{\theta}$ in the next analysis and reinstate it when needed. We also write $p_\theta = p(\mathbf{y};\boldsymbol{\theta})$.

IV. Pearson information matrix

In Section IV-A, we begin with a step-by-step algebraic derivation of $\mathbf{L}$ in (3) which will define the Pearson information matrix. As explained there, the PIM generalizes the results in [5]–[7] and coincides with a bound recently derived in [10] (by comparison this lecture notes provides a simple textbook-style derivation of the bound as well as further connections). In Section IV-B we go on to provide an information-theoretic connection between the PIM and misspecified data distributions using the principle of maximum entropy [11]. Then we study the behaviour of PIM when M increases in Section IV-C. Finally, in Section V, we establish a relation between the PIM and generalized method of moments that is analogous to the relation between the FIM and the maximum likelihood method. The presented results enable a tractable analysis of a wider class of data models that satisfy certain moment constraints.

A. Algebraic derivation

Consider a linear combination of the centered statistics $\mathbf{z}-\boldsymbol{\mu}$:

$$\mathbf{W}^\top(\mathbf{z}-\boldsymbol{\mu}),$$

where $\mathbf{W}^\top \in \mathbb{R}^{n\times M}$ denotes a linear combiner matrix. This vector has zero-mean similar to the score function, cf. (2). We construct the following matrix

$$\mathrm{E}\left[\begin{matrix}\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}} \\ \mathbf{W}^\top(\mathbf{z}-\boldsymbol{\mu})\end{matrix}\right]\left[\begin{matrix}\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}} \\ \mathbf{W}^\top(\mathbf{z}-\boldsymbol{\mu})\end{matrix}\right]^\top \succeq \mathbf{0}. \tag{5}$$

Under regularity conditions that allow the interchanging of integral and derivative operations [2], the following identity holds:

$$\begin{aligned}
&\mathrm{E}\left[\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}}(\mathbf{z}-\boldsymbol{\mu})^\top\right] \\
&= \int \frac{1}{p_\theta}\frac{\partial p_\theta}{\partial \boldsymbol{\theta}}(\mathbf{z}-\boldsymbol{\mu})^\top p_\theta d\mathbf{y} \\
&= \int \left(\frac{\partial p_\theta(\mathbf{z}-\boldsymbol{\mu})^\top}{\partial \boldsymbol{\theta}} + p_\theta \frac{\partial \boldsymbol{\mu}^\top}{\partial \boldsymbol{\theta}}\right) d\mathbf{y} \\
&= \frac{\partial}{\partial \boldsymbol{\theta}} \underbrace{\mathrm{E}\left[(\mathbf{z}-\boldsymbol{\mu})^\top\right]}_{=\mathbf{0}} + \underbrace{\int p_\theta \, d\mathbf{y}}_{=1}\frac{\partial \boldsymbol{\mu}^\top}{\partial \boldsymbol{\theta}} \\
&= \mathbf{D}^\top,
\end{aligned} \tag{6}$$

where

$$\mathbf{D}^\top = \frac{\partial \boldsymbol{\mu}^\top}{\partial \boldsymbol{\theta}} \in \mathbb{R}^{n\times M} \tag{7}$$

is the gradient of the mean vector. Using (4) and (6), the matrix in (5) can be expressed as

$$\begin{bmatrix}\mathbf{J} & \mathbf{D}^\top\mathbf{W} \\ \mathbf{W}^\top\mathbf{D} & \mathbf{W}^\top\boldsymbol{\Sigma}\mathbf{W}\end{bmatrix} \succeq \mathbf{0}. \tag{8}$$

It follows from the Schur complement of the lower-right block of (8) that

$$\mathbf{0} \preceq \mathbf{D}^\top\mathbf{W}(\mathbf{W}^\top\boldsymbol{\Sigma}\mathbf{W})^{-1}\mathbf{W}^\top\mathbf{D} \preceq \mathbf{J}, \tag{9}$$

assuming that $\mathbf{W}^\top\boldsymbol{\Sigma}\mathbf{W}$ has full rank [12]. Equation (9) yields a nonnegative lower bound on the FIM that is dependent on the choice of the linear combiner $\mathbf{W}$.

The tightest lower bound (9) is found by solving the problem

$$\max_{\mathbf{W}} \ \mathbf{D}^\top\mathbf{W}(\mathbf{W}^\top\boldsymbol{\Sigma}\mathbf{W})^{-1}\mathbf{W}^\top\mathbf{D}. \tag{10}$$

The combiner that produces the tightest bound is $\mathbf{W}_\star = \boldsymbol{\Sigma}^{-1}\mathbf{D}$. To show this, begin by constructing the following positive semidefinite matrix,

$$\begin{aligned}
&\begin{bmatrix}\mathbf{D}^\top\boldsymbol{\Sigma}^{-1}\mathbf{D} & \mathbf{D}^\top\mathbf{W} \\ \mathbf{W}^\top\mathbf{D} & \mathbf{W}^\top\boldsymbol{\Sigma}\mathbf{W}\end{bmatrix} \\
&= \begin{bmatrix}\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}^\top\end{bmatrix}\begin{bmatrix}\mathbf{D}^\top\boldsymbol{\Sigma}^{-1}\mathbf{D} & \mathbf{D}^\top \\ \mathbf{D} & \boldsymbol{\Sigma}\end{bmatrix}\begin{bmatrix}\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}\end{bmatrix} \\
&= \begin{bmatrix}\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}^\top\end{bmatrix}\begin{bmatrix}\mathbf{D}^\top\boldsymbol{\Sigma}^{-1/2} \\ \boldsymbol{\Sigma}^{1/2}\end{bmatrix}\begin{bmatrix}\mathbf{D}^\top\boldsymbol{\Sigma}^{-1/2} \\ \boldsymbol{\Sigma}^{1/2}\end{bmatrix}^\top\begin{bmatrix}\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}\end{bmatrix} \succeq \mathbf{0}.
\end{aligned} \tag{11}$$

Using the Schur complement of the lower-right block of (11) we obtain the upper bound

$$\mathbf{D}^\top\mathbf{W}(\mathbf{W}^\top\boldsymbol{\Sigma}\mathbf{W})^{-1}\mathbf{W}^\top\mathbf{D} \preceq \mathbf{D}^\top\boldsymbol{\Sigma}^{-1}\mathbf{D}, \tag{12}$$

which is clearly attained at $\mathbf{W}_\star = \boldsymbol{\Sigma}^{-1}\mathbf{D}$.

In conclusion, using (9) and (12), we have proved the following theorem.

Theorem 1. *The optimal lower bound (in the class of bounds considered) is*

$$\boxed{\mathbf{L} \triangleq \mathbf{D}^\top\boldsymbol{\Sigma}^{-1}\mathbf{D} \preceq \mathbf{J},} \tag{13}$$

where $\boldsymbol{\Sigma}$ and $\mathbf{D}$ are either obtained analytically or computed numerically. We call $\mathbf{L} \succeq \mathbf{0}$ the Pearson information matrix for reasons explained above.

Remark: Suppose $\mathbf{y} \sim p_\theta$ can be modeled as

$$\mathbf{y} = \boldsymbol{\mu}(\boldsymbol{\theta}) + \mathbf{w} \in \mathbb{R}^N,$$

where $\mathbf{w}$ is a zero-mean random variable. Let $\mathbf{z}(\mathbf{y}) = \mathbf{y}$. Then the corresponding PIM coincides with the FIM bounds in [5], [6] and in [7], when the covariance matrix is fixed, $\boldsymbol{\Sigma}$, and variable, $\boldsymbol{\Sigma}(\boldsymbol{\theta})$, respectively. The above algebraic derivation of the PIM provides, moreover, a simple textbook-like proof of the optimized FIM bound in [10] (which also contains an illustrative example consisting of a nonlinear amplification device).

B. Connection to misspecified data distributions

We now relate $\mathbf{L}$ to certain properties of misspecified data distributions using the principle of maximum entropy. Instead of the unknown or intractable distribution p_θ, we will use an alternative statistical model, denoted $p_\star$, along with the following identity:

$$\ln p_\theta = \ln p_\star + \ln \delta_\star,$$

which holds for any choice of $p_\star$, where $\delta_\star = \frac{p_\theta}{p_\star}$.

The uncertainty of the data $\mathbf{y}$ is quantified by the (differential) entropy which can be decomposed as

$$H(p_\theta) \triangleq -\mathrm{E}[\ln p_\theta] = -\mathrm{E}[\ln p_\star] - \Delta(p_\theta||p_\star),$$

where $\Delta(p_\theta||p_\star) = \mathrm{E}[\ln \delta_\star] \geq 0$ is the divergence of $p_\star$ from the unknown distribution p_θ [13], [14]. We decompose the score function into

$$\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}} = \frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}} + \frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}}, \tag{14}$$

where the terms correspond to a misspecified score and a divergence score, respectively. The misspecified information matrix is defined as

$$\mathbf{J}_\star \triangleq \mathrm{E}\left[\frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}}\frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}}^\top\right] \succeq \mathbf{0}, \tag{15}$$

where the expectation is taken with respect to p_θ.

Lemma 1. *A general lower bound on $\mathbf{J}$ is*

$$\mathbf{J}_\star + \widetilde{\mathbf{J}} \preceq \mathbf{J}, \tag{16}$$

where

$$\widetilde{\mathbf{J}} = \mathrm{E}\left[\frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}} \frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}}^\top\right] + \mathrm{E}\left[\frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}} \frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}}^\top\right]. \quad (17)$$

Proof: Inserting (14) into (1), we obtain the following decomposition

$$\mathbf{J} = \mathbf{J}_\star + \widetilde{\mathbf{J}} + \mathrm{E}\left[\frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}} \frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}}^\top\right] \quad (18)$$

and the result follows immediately. ∎

We are concerned with misspecified data models $p_\star$ that satisfy the given constraint $\mathrm{E}[\mathbf{z}] = \boldsymbol{\mu}$. That is, distributions that satisfy

$$\int \mathbf{z} p_\star \, d\mathbf{y} = \boldsymbol{\mu}. \quad (19)$$

In particular, we let $p_\star$ correspond to the maximum uncertainty of $\mathbf{y}$. The distribution with the maximum (differential) entropy is known to be

$$p_\star \equiv \underset{p' \in \mathcal{P}}{\arg\max} \; -\mathrm{E}'[\ln p'] = \exp(\boldsymbol{\lambda}^\top \mathbf{z} - \lambda_0),$$

where $\mathcal{P}$ is the set of valid probability distributions for $\mathbf{y}$ that satisfy (19) and $\lambda_0, \boldsymbol{\lambda}$ are multipliers that are chosen to satisfy the constraint (assuming that the problem is feasible) [14]. For completeness, we prove this result by noting that the following upper bound holds for any p':

$$\begin{aligned} H(p') &= -\mathrm{E}'[\ln p'] \\ &\leq -\mathrm{E}'[\ln p_\star] \\ &= -\int (\boldsymbol{\lambda}^\top \mathbf{z} - \lambda_0) p' \, d\mathbf{y} \\ &= -\int (\boldsymbol{\lambda}^\top \mathbf{z} - \lambda_0) p_\star \, d\mathbf{y} \\ &= H(p_\star). \end{aligned}$$

The equality in the penultimate line follows since both p' and $p_\star$ satisfy the constraint (19). The maximum entropy distribution therefore belongs to the exponential family, that is,

$$p_\star(\mathbf{y}; \boldsymbol{\theta}) = \exp\left(\boldsymbol{\lambda}^\top(\boldsymbol{\theta}) \mathbf{z}(\mathbf{y}) - \lambda_0(\boldsymbol{\theta})\right), \quad (20)$$

where

$$\lambda_0(\boldsymbol{\theta}) = \ln \int \exp\left(\boldsymbol{\lambda}^\top(\boldsymbol{\theta}) \mathbf{z}(\mathbf{y})\right) d\mathbf{y}$$

is a normalizing constant.

Lemma 2. *Using the maximum entropy distribution* (20)*, the bound in* (16) *is given by*

$$\mathbf{J}_\star + \widetilde{\mathbf{J}} = \mathbf{L} - \left(\boldsymbol{\Sigma}^{-1}\mathbf{D} - \frac{\partial \boldsymbol{\lambda}}{\partial \boldsymbol{\theta}}\right)^\top \boldsymbol{\Sigma} \left(\boldsymbol{\Sigma}^{-1}\mathbf{D} - \frac{\partial \boldsymbol{\lambda}}{\partial \boldsymbol{\theta}}\right) \quad (21)$$

Proof: For (20), we have that

$$\begin{aligned} \frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}} &= \frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} \mathbf{z} - \frac{\partial \lambda_0}{\partial \boldsymbol{\theta}} \\ &= \frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} (\mathbf{z} - \boldsymbol{\mu}) \end{aligned} \quad (22)$$

where $\frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}}$ is $n \times M$. The second line follows from the fact that

$$\begin{aligned} \frac{\partial \lambda_0}{\partial \boldsymbol{\theta}} &= \frac{1}{\int \exp(\boldsymbol{\lambda}^\top \mathbf{z}) d\mathbf{y}'} \int \frac{\partial \exp(\boldsymbol{\lambda}^\top \mathbf{z})}{\partial \boldsymbol{\theta}} \, d\mathbf{y} \\ &= \int \frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} \mathbf{z} \left(\frac{\exp(\boldsymbol{\lambda}^\top \mathbf{z})}{\int \exp(\boldsymbol{\lambda}^\top \mathbf{z}) d\mathbf{y}'}\right) d\mathbf{y} \\ &= \int \frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} \mathbf{z} p_\star \, d\mathbf{y} \\ &= \int \frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} \mathbf{z} p_\theta \, d\mathbf{y} \\ &= \frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} \boldsymbol{\mu}. \end{aligned} \quad (23)$$

Next, from the proof of (16) we have that

$$\mathbf{J} - (\mathbf{J}_\star + \widetilde{\mathbf{J}}) = \mathrm{E}\left[\frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}} \frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}}^\top\right]. \quad (24)$$

Using (14), the divergence score can be written as the sum of two random vectors,

$$\begin{aligned} \frac{\partial \ln \delta_\star}{\partial \boldsymbol{\theta}} = &\left(\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}} - \mathbf{D}^\top \boldsymbol{\Sigma}^{-1}(\mathbf{z} - \boldsymbol{\mu})\right) \\ &+ \left(\mathbf{D}^\top \boldsymbol{\Sigma}^{-1}(\mathbf{z} - \boldsymbol{\mu}) - \frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}}\right). \end{aligned} \quad (25)$$

For the maximum entropy distribution, these two random vectors are orthogonal, i.e.,

$$\begin{aligned} &\mathrm{E}\Big[\left(\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}} - \mathbf{D}^\top \boldsymbol{\Sigma}^{-1}(\mathbf{z} - \boldsymbol{\mu})\right) \\ &\quad \times \left(\mathbf{D}^\top \boldsymbol{\Sigma}^{-1}(\mathbf{z} - \boldsymbol{\mu}) - \frac{\partial \ln p_\star}{\partial \boldsymbol{\theta}}\right)^\top\Big] \\ &= \left(\mathbf{D}^\top - \mathbf{D}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\Sigma}\right) \left(\boldsymbol{\Sigma}^{-1}\mathbf{D} - \frac{\partial \boldsymbol{\lambda}}{\partial \boldsymbol{\theta}}\right) \\ &= \mathbf{0}, \end{aligned}$$

where the equality follows from (22) and (6). Finally, by inserting (25) into (24), the right hand side of (24) equals

$$\mathbf{J} - \mathbf{D}^\top \boldsymbol{\Sigma}^{-1} \mathbf{D} + \left(\boldsymbol{\Sigma}^{-1}\mathbf{D} - \frac{\partial \boldsymbol{\lambda}}{\partial \boldsymbol{\theta}}\right)^\top \boldsymbol{\Sigma} \left(\boldsymbol{\Sigma}^{-1}\mathbf{D} - \frac{\partial \boldsymbol{\lambda}}{\partial \boldsymbol{\theta}}\right)$$

and result (21) follows. ∎

Theorem 2. *The tightest bound* (16) *is the PIM:*

$$\boxed{\mathbf{J}_\star + \widetilde{\mathbf{J}} \preceq \mathbf{L} \preceq \mathbf{J},} \quad (26)$$

and the corresponding misspecified information matrix is $\mathbf{J}_\star = \mathbf{L}$.

Proof: It follows from (22) that for maximum entropy distributions, the misspecified information matrix is

$$\mathbf{J}_\star = \frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma} \frac{\partial \boldsymbol{\lambda}}{\partial \boldsymbol{\theta}}.$$

Furthermore, in (21) it is readily seen that the tightest bound is attained for

$$\frac{\partial \boldsymbol{\lambda}^\top}{\partial \boldsymbol{\theta}} = \mathbf{D}^\top \boldsymbol{\Sigma}^{-1}. \quad (27)$$

Therefore (27) leads to $\mathbf{J}_\star = \mathbf{L}$ and $\widetilde{\mathbf{J}} = \mathbf{0}$ in (16). ∎

C. The PIM increases as M increases

The vector $\mathbf{z}$ employs M statistics, and to stress that we write $\mathbf{L}_M = \mathbf{D}_M^\top \mathbf{\Sigma}_M^{-1} \mathbf{D}_M$.

Theorem 3. *Including more statistics in* $\mathbf{z}$ *can never worsen the bound, i.e.,*

$$\mathbf{0} \preceq \mathbf{L}_n \preceq \cdots \preceq \mathbf{L}_M \preceq \mathbf{L}_{M+1} \preceq \mathbf{J}. \tag{28}$$

Proof: Suppose we extend the vector $\mathbf{z}$ with an $(M+1)$th statistic so that we can write

$$\mathbf{\Sigma}_{M+1} = \begin{bmatrix} \mathbf{\Sigma}_M & \mathbf{c} \\ \mathbf{c}^\top & \kappa \end{bmatrix} \quad \text{and} \quad \mathbf{D}_{M+1} = \begin{bmatrix} \mathbf{D}_M \\ \mathbf{d}^\top \end{bmatrix}. \tag{29}$$

Then calculating $\mathbf{\Sigma}_{M+1}^{-1}$ using [15, Lemma A.2], we obtain

$$\begin{aligned} \mathbf{L}_{M+1} &= \begin{bmatrix} \mathbf{D}_M^\top & \mathbf{d} \end{bmatrix} \begin{bmatrix} \mathbf{\Sigma}_M & \mathbf{c} \\ \mathbf{c}^\top & \kappa \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{D}_M \\ \mathbf{d}^\top \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{D}_M^\top & \mathbf{d} \end{bmatrix} \Bigg(\begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} \mathbf{\Sigma}_M^{-1} \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix} \\ &\quad + \frac{1}{\kappa - \mathbf{c}^\top \mathbf{\Sigma}_M^{-1} \mathbf{c}} \begin{bmatrix} -\mathbf{\Sigma}_M^{-1}\mathbf{c} \\ 1 \end{bmatrix} \begin{bmatrix} -\mathbf{c}^\top \mathbf{\Sigma}_M^{-1} & 1 \end{bmatrix} \Bigg) \begin{bmatrix} \mathbf{D}_M \\ \mathbf{d}^\top \end{bmatrix} \\ &= \mathbf{L}_M + \frac{(\mathbf{d} - \mathbf{D}^\top \mathbf{\Sigma}_M^{-1}\mathbf{c})(\mathbf{d} - \mathbf{D}^\top \mathbf{\Sigma}_M^{-1}\mathbf{c})^\top}{\kappa - \mathbf{c}^\top \mathbf{\Sigma}_M^{-1}\mathbf{c}} \\ &\succeq \mathbf{L}_M. \end{aligned}$$

■

An interesting research problem is to study the limit of $\mathbf{L}_M$ as $M \to \infty$. Under what conditions will $\mathbf{L}_M$ converge to $\mathbf{J}$?

V. PIM and Generalized method of moments

An efficient unbiased estimator $\widehat{\boldsymbol{\theta}}$ exists if and only if the following identity holds [3], [4]

$$\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}} = \mathbf{J}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}), \tag{30}$$

which is satisfied only in rare cases. In more general scenarios, the maximum likelihood estimator

$$\widehat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} \; -\ln p(\mathbf{y}; \boldsymbol{\theta}). \tag{31}$$

is (asymptotically) unbiased with (asymptotic) covariance matrix $\mathrm{Cov}[\widehat{\boldsymbol{\theta}}] = \mathbf{J}^{-1}$ under certain regularity conditions. It is thus asymptotically efficient. Given an appropriate initialization point $\widehat{\boldsymbol{\theta}}_0$, (31) can be solved iteratively using the Newton-based scoring method:

$$\widehat{\boldsymbol{\theta}}_{i+1} = \widehat{\boldsymbol{\theta}}_i - \mathbf{J}^{-1} \frac{\partial - \ln p_\theta}{\partial \boldsymbol{\theta}}\Big|_{\boldsymbol{\theta} = \widehat{\boldsymbol{\theta}}_i}. \tag{32}$$

When $\mathbf{J}$ and $\frac{\partial \ln p_\theta}{\partial \boldsymbol{\theta}}$ are unknown, we may use instead the best misspecified information matrix (15) and score (22), in the sense of Theorem 2. Then the so-obtained scoring method is related to a method that minimizes a certain class of cost functions $V(\boldsymbol{\theta})$ which we characterize in what follows. Consider evaluating the misspecified variables (15) and (22) obtained using (27) and a consistent estimate $\widehat{\mathbf{\Sigma}}$ instead of $\mathbf{\Sigma}$. Then the corresponding scoring method, analogous to (32), can be expressed as

$$\widehat{\boldsymbol{\theta}}_{i+1} = \widehat{\boldsymbol{\theta}}_i - \widehat{\mathbf{J}}_\star^{-1} \frac{\partial - \ln \widehat{p}_\star}{\partial \boldsymbol{\theta}}\Big|_{\boldsymbol{\theta} = \widehat{\boldsymbol{\theta}}_i}., \tag{33}$$

where we define (according to the above discussion about (15), (22) and (27))

$$\begin{aligned} \frac{\partial \ln \widehat{p}_\star}{\partial \boldsymbol{\theta}} &= \mathbf{D}^\top(\boldsymbol{\theta}) \widehat{\mathbf{\Sigma}}^{-1} (\mathbf{z} - \boldsymbol{\mu}(\boldsymbol{\theta})) \\ \widehat{\mathbf{J}}_\star &= \mathbf{D}^\top(\boldsymbol{\theta}) \widehat{\mathbf{\Sigma}}^{-1} \mathbf{D}(\boldsymbol{\theta}). \end{aligned} \tag{34}$$

One can verify that (33) is a scoring method for solving the following problem

$$\widehat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}} \; \underbrace{\frac{1}{2} (\mathbf{z} - \boldsymbol{\mu}(\boldsymbol{\theta}))^\top \widehat{\mathbf{\Sigma}}^{-1} (\mathbf{z} - \boldsymbol{\mu}(\boldsymbol{\theta}))}_{\triangleq V(\boldsymbol{\theta})}, \tag{35}$$

by noting that $\partial_\theta V(\boldsymbol{\theta}) = -\frac{\partial \ln \widehat{p}_\star}{\partial \boldsymbol{\theta}}$ and that $\widehat{\mathbf{J}}_\star$ is an estimate of the Hessian $\partial_\theta^2 V(\boldsymbol{\theta})$. Eq. (35) is recognized as a generalized method of moments, using an asymptotically optimal weight matrix $\widehat{\mathbf{\Sigma}}^{-1}$ [9], [15].

The cost function $V(\boldsymbol{\theta})$ can be characterized around its minimum, as follows

$$\mathbf{0} = \partial_\theta V(\widehat{\boldsymbol{\theta}}) \simeq \partial_\theta V(\boldsymbol{\theta}) + \partial_\theta^2 V(\boldsymbol{\theta})(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}), \tag{36}$$

where the right-hand side is a Taylor expansion. Using properties of (34) in (36), we can solve for $\widehat{\boldsymbol{\theta}}$ and obtain the following approximation

$$\widehat{\boldsymbol{\theta}} \simeq \boldsymbol{\theta} + (\mathbf{D}^\top \widehat{\mathbf{\Sigma}}^{-1} \mathbf{D})^{-1} \mathbf{D}^\top \widehat{\mathbf{\Sigma}}^{-1} (\mathbf{z} - \boldsymbol{\mu}).$$

Since the unknown distribution p_θ satisfies (19), it follows that $\mathrm{E}[\widehat{\boldsymbol{\theta}}] \simeq \boldsymbol{\theta}$ and

$$\mathrm{Cov}[\widehat{\boldsymbol{\theta}}] \simeq (\mathbf{D}^\top \mathbf{\Sigma}^{-1} \mathbf{D})^{-1} = \mathbf{L}^{-1} \succeq \mathbf{J}^{-1}.$$

The above expressions hold asymptotically as the number of samples N in $\mathbf{y}$ increases [15].

In summary, using a scoring method analogous to (32) leads to the generalized method of moments (35) with asymptotic covariance given by the inverse PIM.

VI. Conclusions

We have provided a direct, algebraic derivation of a tractable lower bound on the Fisher information matrix which we called the Pearson information matrix (for reasons explained above). Furthermore, we presented an information-theoretic link between the PIM and misspecified data distributions as well as a connection to the generalized method of moments.

References

[1] C. Rao, "Information and accuracy attainable in the estimation of statistical parameters," *Bulletin of the Calcutta Mathematical Society*, vol. 37, no. 3, pp. 81–91, 1945.

[2] H. Cramér, "A contribution to the theory of statistical estimation," *Scandinavian Actuarial Journal*, vol. 1946, no. 1, pp. 85–94, 1946.

[3] H. Van Trees and K. Bell, *Detection Estimation and Modulation Theory, Pt.I.* Detection Estimation and Modulation Theory, Wiley, second ed., 2013 [1968].

[4] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation theory*. Prentice Hall, 1993.

[5] P. Stoica and P. Babu, "The Gaussian data assumption leads to the largest Cramér-Rao bound," *IEEE Signal Processing Mag.*, vol. 28, no. 3, pp. 132–133, 2011.

[6] S. Park, E. Serpedin, and K. Qaraqe, "Gaussian assumption: The least favorable but the most useful," *IEEE Signal Processing Mag.*, vol. 30, no. 3, pp. 183–186, 2013.

[7] M. Stein, A. Mezghani, and J. Nossek, "A lower bound for the Fisher information measure," *IEEE Signal Processing Letters*, vol. 21, pp. 796–799, July 2014.
[8] K. Pearson, "Contributions to the mathematical theory of evolution," *Philosophical Transactions of the Royal Society of London. A*, vol. 185, pp. 71–110, 1894.
[9] L. P. Hansen, "Large sample properties of generalized method of moments estimators," *Econometrica: Journal of the Econometric Society*, pp. 1029–1054, 1982.
[10] M. Stein, J. A. Nossek, and K. Barbé, "Fisher information bounds with applications in nonlinear learning, compression and inference," *arXiv preprint arXiv:1512.03473*, 2015.
[11] E. T. Jaynes, "Information theory and statistical mechanics," *Physical review*, vol. 106, no. 4, p. 620, 1957.
[12] R. Horn and C. Johnson, *Matrix Analysis*. Cambridge University Press, 1990.
[13] S. Kullback, *Information Theory and Statistics*. Dover, 1997 [1959].
[14] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, 2006.
[15] T. Söderström and P. Stoica, *System identification*. Prentice-Hall, Inc., 1988.